Multiple Sensor Platform Coordination Using Stigmergy

March 1998

Barrett F. Richey
BTG, Inc.
3877 Fairfax Ridge Road, 3C
Fairfax, VA, 22030-7448
brichey@btg.com

ABSTRACT

The current investigation explores the use of a biological technique, known as Stigmergy, to coordinate semi-autonomous surveillance platforms during search operations such as in the case of Unmanned Aerial Vehicles (UAVs). Stigmergy was recognized and named by French biologist P.P. Grass`e in 1959 while studying nest building of termites. He observed that indirect coordination among termites was accomplished through sensing and modification of their environment. The chemicals secreted by each termite during nest building affected the building actions of neighboring termites, resulting in a coordinated building strategy. Stigmergy is not limited to termites. Some species of ants use stigmergy for trail recruitment, where the interactions among foragers are mediated by pheromones they leave as a trail.[Beckers *et al.* 1994] A series of experiments investigates the use of stigmergy to coordinate multiple sensor platforms in a search operation for a stationary, ground target. The purpose of these experiments is to compare the search time of the stigmergic strategy to an independent, uncoordinated strategy and to a mechanically coordinated strategy.

1.0 INTRODUCTION

From Desert Storm accounts it is apparent that hunting and destroying mobile missile launchers, carrying weapons of mass destruction, was difficult to say the least. More than a year after its conclusion, the coalition Air Component Commander during Desert Storm, Lt. Gen. Horner, delivered the keynote speech for the Theater Missile Defense Symposium at the U.S. State Department. He talked about the valiant effort made to stop Iraq from launching SCUDs directed at Israel. The significant presence of Combat Air Patrol (CAP) fighters suppressed the launches, but he believed that they did not destroy a single Tractor Erector Launcher (TEL). This failure was repeated four and a half years later at Roving Sands ’95 which achieved very limited success. The improvements made during the four and a half years in command and control, computer connectivity and sensors did not make much, if any, difference.

The use of UAVs has proliferated in an attempt to reduce sensor platform costs, cover more of the battlefield, and get the sensors closer to the objects to aid in target identification. Much

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (FROM - TO)
01-03-1998	Conference Proceedings	xx-xx-1998 to xx-xx-1998

4. TITLE AND SUBTITLE
Multiple Sensor Platform Coordination Using Stigmergy
Unclassified

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Richey, Barrett F. ;

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME AND ADDRESS
BTG, Inc.
3877 Fiarfax Ridge Road, 3C
Fairfax, VA22030-7448

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME AND ADDRESS
Director, CECOM RDEC
Night Vision and Electronic Sensors Directorate, Security Team
10221 Burbeck Road
Ft. Belvoir, VA22060-5806

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
APUBLIC RELEASE
,

13. SUPPLEMENTARY NOTES
See Also ADM201041, 1998 IRIS Proceedings on CD-ROM.

14. ABSTRACT
The current investigation explores the use of a biological technique, known as Stigmergy, to coordinate semi-autonomous surveillance platforms during search operations such as in the case of Unmanned Aerial Vehicles (UAVs). Stigmergy was recognized and named by French biologist P.P. Grass`e in 1959 while studying nest building of termites. He observed that indirect coordination among termites was accomplished through sensing and modification of their environment. The chemicals secreted by each termite during nest building affected the building actions of neighboring termites, resulting in a coordinated building strategy. Stigmergy is not limited to termites. Some species of ants use stigmergy for trail recruitment, where the interactions among foragers are mediated by pheromones they leave as a trail.[Beckers et al. 1994] A series of experiments investigates the use of stigmergy to coordinate multiple sensor platforms in a search operation for a stationary, ground target. The purpose of these experiments is to compare the search time of the stigmergic strategy to an independent, uncoordinated strategy and to a mechanically coordinated strategy.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:

a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19. NAME OF RESPONSIBLE PERSON
Unclassified	Unclassified	Unclassified	Public Release	10	Fenster, Lynn lfenster@dtic.mil

19b. TELEPHONE NUMBER
International Area Code
Area Code Telephone Number
703767-9007
DSN
427-9007

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39.18

research and development has gone into UAVs and sensor technology, but little research has been done to devise ways to coordinate the battlefield sensors. How can numerous sensor platforms be made to operate as a coordinated unit even as they perform their independent actions? To coordinate sensor platforms this paper investigates the use of stigmergy, which enables ants to coordinate their efforts even as each ant performs its autonomous task.

The concept explored here is based in large measure on the work of [Beckers *et al.* 1994]. They studied the use of stigmergy to coordinate miniature robots that pushed small wooden pucks into a single pile, demonstrating that simple, less expensive robots could perform coordinated activities that might otherwise require a single, much more expensive robot. The robots actively sensed and modified their environment (the pucks) in their efforts to push the pucks into a pile. Briefly, three to five robots were placed in a bounded area where 81 small pucks were evenly distributed. The goal was to push the pucks into a single pile, even though each robot was designed to push no more than three pucks at any one time. If a forth puck was obtained, the robot would back up, releasing the pucks, and turn a random angle before moving forward again. The robots' actions would cause pucks to be knocked free of piles with more than three pucks, making the loose pucks available to push. The action of one robot, in pushing pucks and leaving small piles of pucks, affected the action of the other robots. The robots successfully pushed all of the pucks into a single pile, coordinating their efforts with stigmergy. The authors' use of stigmergy stems from research on ants and other social insects that sense and modify their environment by excreting chemicals which affects the actions of others in the colony.

> "[Stigmergy] also occurs in cooperative foraging strategies such as trail recruitment in ants, where the interactions between foragers are mediated by pheromones put on the ground in quantities determined by the local conditions of the environment. For example, trail recruiting ant species are able to select and preferentially exploit the richest food source in the neighborhood (Pasteels, Deneubourg and Goss 1987; Beckers *et al*. 1990) or the shortest path between the nest and a food source (Beckers, Deneubourg and Goss 1992). This strategy takes advantage of the characteristics of the trail-laying and trail-following mechanisms of the ants in combination with their essentially probabilistic behavior: the probability that an ant follows a trail is a non-linear function of the trail's pheromone concentration, and the probability that an ant lays a pheromone spot depends on the characteristics of the recently-encountered food source and the environment (Beckers, Deneubourg and Goss 1994). When a trail between a single food source and the nest is first established, its pheromone concentration is low, and a high proportion of ants lose the path before reaching the food or the nest. As more journeys are made along the trail, the pheromone concentration increases progressively and so does the accuracy of trail following; finally the majority of the foragers will successfully use that trail."[Beckers *et al*. 1994]

Since UAVs are, in effect, robots that sense their environment, the use of stigmergy to coordinate their activities may have potential. Unlike ants and the miniature robots, UAVs are not able to modify their physical environment, so the stigmergic technique described in this paper makes use of the concept of a virtual environment to act as the medium for indirect communications among the sensor platforms. This virtual environment is a computer model that represents the physical battlefield as a 2-dimensional matrix, where each cell in the matrix represents a defined square on the battlefield.

The matrix is initialized with an area delimitation product, such as IBIS or GALE suitability data, for the target type and target state (e.g., SCUD TEL in hide state) being sought. The suitability data aids in directing sensor platforms to areas more suited for the target. If the area delimitation product is accurate in quantifying the target's preference to occupy a given cell, surveying more suitable cells should increase the chances of locating the target. Although the matrix is initialized with *a priori* terrain suitability values for a specific target in a specific state, the matrix is modified during the search operation based on the observations of the sensor platforms, transforming it into a likelihood matrix. If a sensor that is appropriate for finding the specific target type in the specific state surveys an area and finds a target, then the likelihood of the target being in that area is increased. The target detection is not a certain fact, since sensors have the potential to generate false alarms. In contrast, if a sensor that is appropriate for finding the specific target type in the specific state surveys an area and finds no target, then the likelihood of the target being in that area is decreased. Figure 1 shows how this matrix would be integrated into the UAV ground station operations.

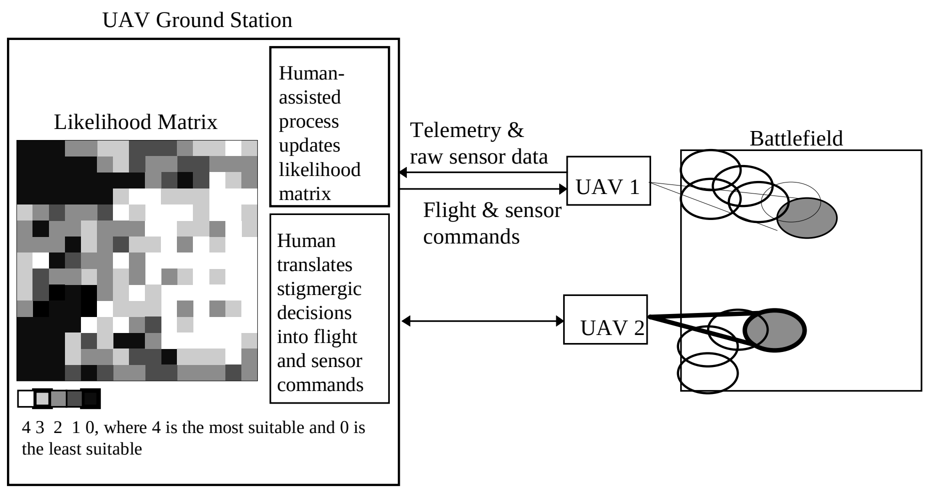

Figure 1: The battlefield, represented on the right, is being surveyed by two UAVs. The filled ellipses represent the areas being surveyed. The unfilled ellipses represent the areas already surveyed. The telemetry and sensor data are processed back at the ground station, updating the likelihood matrix. The likelihood matrix is modified with the negative information as shown by the black squares on the left that correspond to the areas already surveyed on the right. Both UAVs move toward the center of the battlefield where the likelihood is higher than the upper and lower edges.

Stigmergy makes use of the fact that no targets are detected in a cell. This is referred to as negative information and is incorporated into the likelihood matrix. Sensors can be assigned a probability of detecting a specific target type in a specific target state.[*] When a sensor surveys a portion of the battlefield and finds nothing, the negative information is combined with the likelihood values assigned to the cells that correspond to the surveyed portion of the battlefield. The matrix is transformed by the negative information using the formula below.

$$\Lambda_{ij} = 2\lambda_{ij}(1 - P_D) \qquad \{P_D > 0.5\}$$

[*] Multiple target types, multiple target state, and the full use of negative and positive information from target detection and identification is not investigated in this paper. The experiments simulate a single, hidden, single-state target and sensor platforms carrying a simple, generic sensor.

where λ_{ij} is the current likelihood of the cell in the i^{th} row and j^{th} column;
P_D is the probability of detection by the sensor; and
Λ_{ij} is the resulting likelihood.

If no targets are detected and the P_D is 1.0, then Λ_{ij} is 0.0. In other words, a sensor with a perfect ability to detect a target will render a surveyed area with zero likelihood that a target exists. If no targets are detected and the P_D is 0.5, then Λ_{ij} is unchanged. In other words, a sensor with a 50:50 chance of detecting a target provides no useful negative information. Once the likelihood matrix is updated, the software that controls the UAV's position interrogates the likelihood matrix. The algorithm interrogates the cells surrounding the UAV's current location to determine the direction with the highest measure of likelihood. This algorithm is described in more detail below.

The negative information deposited in the likelihood matrix is analogous to the pheromone excreted by an ant. Where the ant's pheromone trail attracts ants, the negative information trail left by a UAV repels UAVs. All UAVs involved in the stigmergic strategy automatically avoid areas already searched in favor of areas suitable to the target that have not yet been searched. Not only does the negative information deposited in the likelihood matrix repel others UAVs, it also repels the UAV that provided the information. This fact is exploited by the first two experiments as described in sections 2.1 and 2.2. It should be noted here that the formula above does not take into account a temporal component. Just as an ant's pheromone trail evaporates over time, so should the likelihood deposited by a sensor platform. However, the target model used in these experiments is stationary, eliminating the need for evaporation. Future work will investigate this aspect of stigmergy.

A benefit of this stigmergic approach is that platforms can be added to and removed from the search strategy without changing the individual strategies of the participating sensor platforms or the overall strategy. Just as a few ants, more or less, have no affect on each ant's individual strategy and have no affect on their ability to find and return food to the colony, the stigmergic approach requires no modification to the central strategy.

2.0 EXPERIMENTS

This investigation is comprised of four experiments designed to explore and compare three search strategies. The first strategy is a basic uncoordinated strategy that is used to establish the baseline for the comparison of the other search strategies. The second strategy employs a simple, Area of Responsibility (AOR) approach to coordinate multiple sensor platforms. The third strategy employs stigmergy to coordinate multiple sensor platforms.

All experiments use the same 1024x1024 likelihood matrix which is initialized with a uniform distribution of integer, suitability values ranging from zero to four, where a larger number represents a higher suitability than a smaller number. All experiments employ the same sensor platform model that surveys the cell in which it resides. The model moves to and surveys one cell for each simulation time step. The sensor has a probability of detecting the target (P_D) of 1.0. The algorithm to determine a sensor platform's next move interrogates the likelihood matrix in eight directions corresponding to up, down, left, right, upper left, lower right, upper right, and lower left as shown in Figure 2. It can be likened to an ant that "sniffs" the ground to determine the pheromone nearby. In completely independent work the term sniff is also used by [Resnick 1994] to denote a similar activity in his simulated turtles, termites and ants.

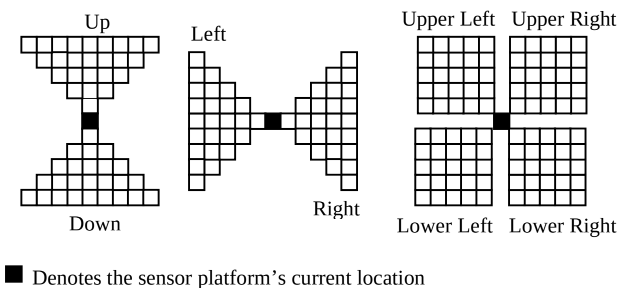

Figure 2: The flight control algorithm (or sniff strategy) interrogates the likelihood matrix in eight directions, looking for the direction that provides the highest likelihood of there being a target. The figure shows all eight directions given a sniff distance of five.

The sniff strategy used throughout these experiments computes the likelihood of finding a target in each of the eight directions. The likelihood score for each direction is computed based on the current sniff distance. The current study uses a standard distance of five. Each direction score is computed by a weighted sum of the cell values in that direction. The total number of cells involved depends on the sniff distance. Examples of the up and upper-right directions are shown in Figure 3. If all direction scores return zero or if more than one return a non-zero tie score, the sniff distance is incremented by five and new scores are computed until a non-zero/non-tie situation occurs. Once the move direction is determined, the sniff distance is reset to five in preparation for the next invocation.

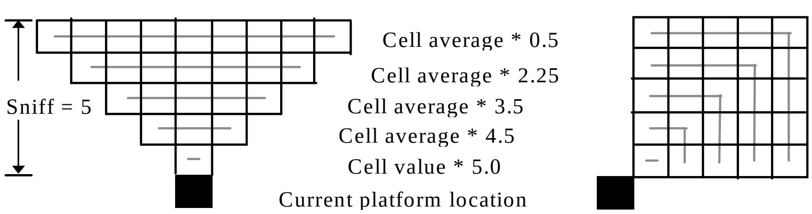

Figure 3: A direction is sniffed by summing the weighted averages at each step. The black square represents the platform's current location. The left side of the figure shows the up direction with a sniff distance of five. The right side of the figure shows the upper-right direction with a sniff distance of five.

The sniff strategy used in these experiments applies a weight to the cell values based on the distance from the sensor platform. As shown in the center of Figure 3 the weights used in these experiments favor cells closer to the platform. Although not investigated in this paper, other sniff strategies may prove to be more effective.

2.1 Uncoordinated Search Strategy Experiment

This first experiment serves as the baseline, computing the average number of time steps it takes for a single sensor platform to find a hidden target in the simulated environment of 1024x1024 cells. It also serves to determine the number of random target locations required to minimize statistical error. Sixty target locations are randomly selected, with a distribution of:

- 24 targets in locations assigned a suitability of 4 (the most suitable),
- 18 in locations assigned a suitability of 3,
- 12 in locations assigned a suitability of 2,
- 6 in locations assigned a suitability of 1 (the least suitable), and
- 0 in locations assigned a suitability of 0 (unsuitable).

Six sensor platforms, starting from the first row and columns that are evenly distributed, search for each target, one target at a time as shown in Figure 4. A target is hidden in the first of the 60 locations. The simulation time is set to zero, and the first of the six platforms searches for the target, using the sniff strategy outlined earlier.

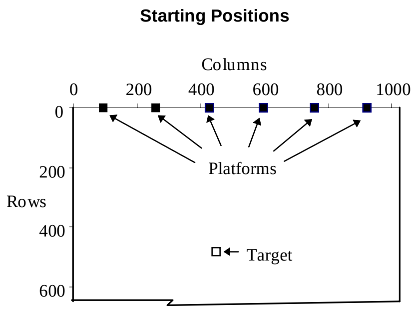

Figure 4: The six platforms, represented as black squares, are evenly distributed across the first row and are ready to begin searching for one of the targets, represented by the white square, in the target sequence

Each sensor platform updates its own copy of the likelihood matrix so that no sharing of negative information is allowed. In this case each sensor platform uses its own likelihood matrix which is initialized with the suitability data generated for these experiments. Each platform deposits its own negative information so that it doesn't survey areas it already surveyed. When the platform finds the target, the current simulation time is recorded. Once the target is located by the first platform, the timer is reset and the second platform begins its search. Once the target is found by all six platforms, the mean of the time steps is recorded. The mean is used to remove any bias due to a platform's starting location. The platforms are reset to their starting locations, and the target is hidden in the next target location in the sequence. Processing continues in this fashion until all six platforms find all 60 targets. The results of the first experiment are shown in Figure 5.

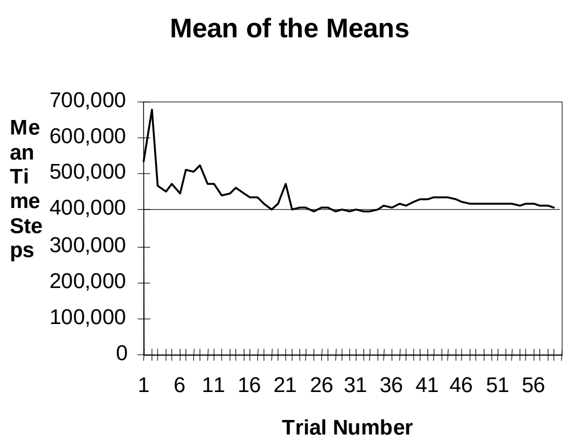

Figure 5: The mean of the mean search times is fairly stable by the 30th target. For the simulated battlefield used in these experiments, the mean, uncoordinated search time is approximately 400,000 time steps.

Figure 5 shows just one random ordering of the trial means. From reviewing many random orderings, the mean number of time steps to locate a target stabilizes at approximately 400,000 steps at or before the 30th trial, thus subsequent experiments use 30 trials. The 400,000 time steps serves as the benchmark for uncoordinated search on the 1024x1024 battlefield model.

2.2 Area of Responsibility (AOR) Search Strategy Experiment

The AOR experiment establishes a second set of numbers to compare with the stigmergic strategy. The AOR search strategy divides the battlefield into evenly sized columns, one column for each sensor platform. For example, when two platforms search for a target, one platform begins at the first row and column 256 and searches from column 1 to column 512 while the other platform begins at the first row and column 768 and searches from column 513 to column 1024 as shown in Figure 6.

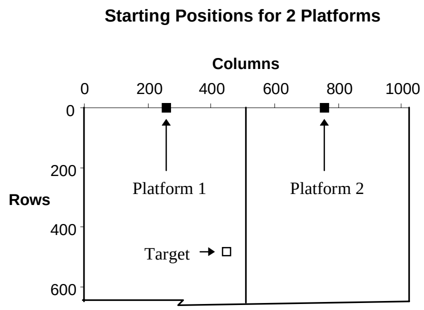

Figure 6: The starting locations for two platforms, using the AOR strategy, are represented by the two, black squares along the top of the battlefield. The area to the left of the dividing line is searched by platform 1, and the area on the right is searched by platform 2.

For the AOR experiment 30 target locations are randomly selected, with a target distribution of

- 12 targets in locations assigned a suitability of 4,
- 9 in locations assigned a suitability of 3,
- 6 in locations assigned a suitability of 2,
- 3 in locations assigned a suitability of 1, and
- 0 in locations assigned a suitability of 0 as shown in Figure 7.

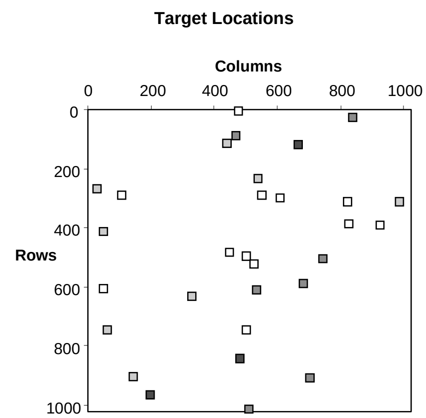

Figure 7: The simulated battlefield is a 2-dimensional array (1024x1024). The filled rectangles identify the target locations. The darker rectangles represent targets hidden in less suitable locations, while the lighter rectangles represent targets hidden in more suitable locations.

In an experiment similar to the first, a target is hidden in the first of the 30 locations and the simulation timer is set to zero. The process begins with two platforms searching for the target. Each platform uses the sniff strategy outlined earlier. Each sensor platform updates its own copy of the likelihood matrix so that no sharing of negative information is allowed. In this case information sharing is not necessary since each platform searches a different, predefined area of the battlefield. Since the platforms are searching for the target in parallel, all platforms survey one cell per time step. When one of the platforms finds the target, searching is halted for that target and the current simulation time is recorded. In this experiment only one platform ever finds the hidden target, since the target can only reside in one AOR. Once the target is found, the platforms are reset to their starting locations; the timer is reset; and the target is hidden in the next target location in the sequence. Processing continues in this fashion until all 30 targets are found. Once the 30 targets are found by the two platforms, processing repeats with 4 platforms, then 8, 16, 32, and finally, 64 platforms. The results of the experiment are shown in Figure 8. The horizontal line at the top of the graph represents the results from the first experiment for purposes of comparison. The graph shows a significant decrease in time steps using the AOR strategy.

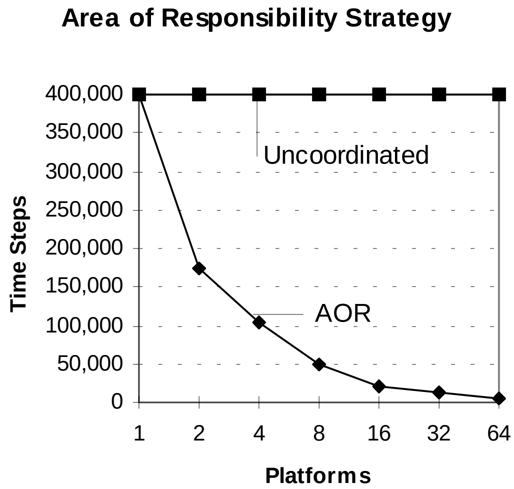

Figure 8: From Figure 5 any number of uncoordinated platforms average approximately 400,000 time steps to find the target regardless of the number of platforms. The AOR strategy shows a steady decrease in time steps as more platforms are added.

As shown in the Figure 8, the doubling of sensor platforms reduces the search time by approximately one half as expected, since the size of the AORs is also reduced by one half.

2.3 Stigmergic Search Strategy Experiments

The first stigmergy experiment employs the same 30 target locations as the AOR experiment above. It also follows the same process whereby two platforms search for the target hidden in the first location. Once the target is found, the current time step is recorded, then the platforms are reset to their starting locations; the timer is reset; and process is repeated to search for a target hidden in the second location. This activity is repeated for all 30 targets. Once all 30 targets are found by the two sensor platforms, processing continues with 4 sensor platforms, then 8, 16, 32, and finally, 64 platforms. In each case the platforms start on the first row and are evenly distributed across the top of the battlefield just as in the AOR experiment. The sensor platforms coordinate their search by sharing a common likelihood matrix as described in the introduction and pictured in Figure 1. The results, shown in Figure 9, are comparable to the AOR experiment. Recall that the use of stigmergy does not require any predefined areas of responsibility based on the number of platforms.

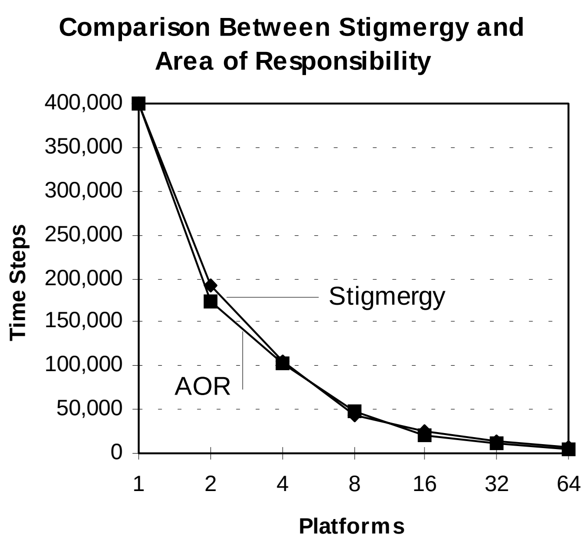

Figure 9: The stigmergic strategy shows a comparable reduction in time steps as platforms are added, but stigmergy does not require a predefined strategy that is dependent on the number of platforms in the search.

The AOR experiment employs a direct strategy of assigning an AOR to each platform and assigning a starting location for each platform. Once the search begins any change in the number of platforms involved in the search must be directly communicated to the remaining platforms to update the strategy. The stigmergic strategy does not require such a change in strategy when a platform is added to or subtracted from the search operation. However, the first stigmergy experiment assigned a starting location for each platform. Placing them evenly along the entry row.

A second stigmergy experiment was conducted which did not rely on any predefined strategy, starting all sensor platforms at the same location (row 1 and column 512). The two experiments yielded comparable results as shown in Figure 10.

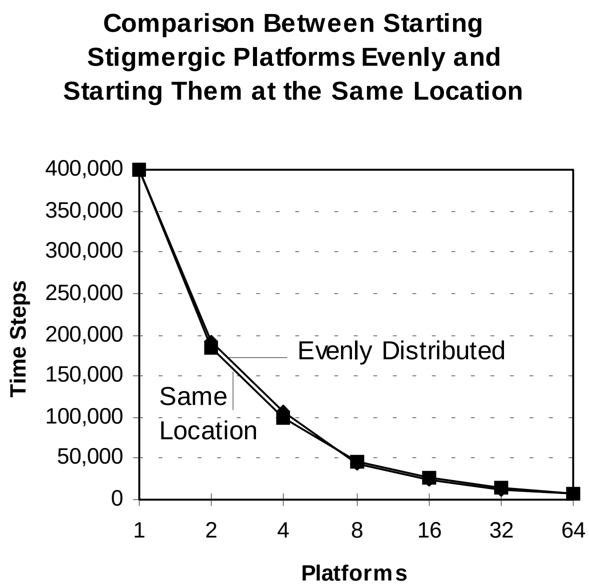

Figure 10: Starting all stigmergic platforms at the same location made little if any difference in their performance.

Since the negative information left behind by stigmergic platforms repels other stigmergic platforms, the platforms, starting at the same location, scattered quickly. The result was nearly identical to the platforms that were evenly distributed across the entry row.

3.0 DISCUSSION

It has been shown experimentally that the use of stigmergy in a semi-autonomous search strategy yielded results similar to a coordinated search strategy based on an assigned Area of Responsibility (AOR). Keep in mind that the stigmergic strategy does not require that AORs be assigned at any time based on the number of platforms. Also, the stigmergic strategy does not require any adjustment if the number of platforms is increased or decreased at any time during the operation. This is important for the following reasons:

1. The initial strategy development, whether AOR or some other direct strategy, becomes more complex as the number of platforms increases.
2. Tactical sensors may be diverted from their assigned strategy. A sensor platform may be grounded, shot down, recalled early, or re-tasked for other purposes, such as Battle Damage Assessment, affecting the initially assigned strategy.
3. Sensor platforms, by design, come on and off station at irregular times. This makes computing a comprehensive strategy very difficult.
4. It is difficult to compute and communicate a change in strategy once the mission has begun.

Although the experiments modeled UAV-like sensor platforms, the stigmergic approach is easily extended to include other platform types. Broad-area surveillance can be added to enhance the UAV search strategy. Broad-area surveillance provides enormous amounts of negative information which currently goes unused. In addition, defense of the group, like search and retrieval of food, is another activity exhibited by social insects.[Beckers *et al.* 1994] With this in mind, ground-to-air threats can be added to the information kept in the shared matrix. The cells surrounding the locations of ground-to-air threats can be designated so that the sniff algorithm prevents the UAVs from entering the threatened areas. The areas not covered by the UAVs can be automatically (stigmergically) covered by assets that can survey the areas without entering them.

4.0 FUTURE DIRECTION

The experiments outlined above modeled simple sensor platforms and a single, stationary target on a simple, white noise battlefield. The positive results shown here encourage plans for more ambitious future work.

The next step is to run experiments on a more robust battlefield model, using actual area delimitation products of different, “interesting” areas of the world. There are many aspects of stigmergy worth investigating, namely:

1. multiple sensor platforms with varying P_Ds,
2. different sniff algorithms,
3. adding and deleting platforms from the strategy,
4. multiple target states and moving between states (e.g., launch, hide, reload),
5. multiple target types, each with their own set of states, and
6. threat avoidance and sensor coverage due to threats.

The eventual goal is to integrate stigmergy into working UAVs and other sensor assets. Even if sensors are not fully integrated into the stigmergic strategy, the negative information supplied by non-participants can be passed to the UAV ground station and used to modify the likelihood matrix, thereby reducing the search space and time of the participating sensor platforms.

Acknowledgments

Thanks go out to Randall Fuerst, Senior Vice President, for supporting this effort, to my friend and colleague Mike Kenworthy for encouragement and critique, and to my wife Malou for her unwavering love and support.

References

[Beckers *et al.* 1994] Beckers, R., Holland, O.E., Deneubourg, J.L., “From Local Actions to Global Tasks: Stigmergy and Collective Robotics," *Artificial Life IV Proceedings*, MIT Press, Cambridge, MA, 1994

[Resnick 1994] Resnick, M., *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*, MIT Press, Cambridge, MA, 1994, page 33